

Geoff Huston
November 2017

Thanks Google!

When writing this blog I am never sure who reads it and how they take it.

Back in August (<http://www.potaroo.net/ispcol/2017-08/xtn-hdrs.html>) I wrote up an examination of the behaviour of IPv6 and packet fragmentation with an observation relating to Google's Public DNS Service. I used two example cases of queries for Google's Public DNS Service to resolve:

In the first case the DNS response is 1,190 octets in length, and in the second case the response is 1,346 octets in length. The DNS server is a IPv6-only server, and the underlying host of this name server is configured with a local maximum packet size of 1,280 octets. This means that in the first case the response being sent to the Google resolver is a single unfragmented IPv6 UDP packet, and the second case the response is broken into two fragmented IPv6 UDP packets. And it is this single change that triggers the Google Public DNS Server to provide the intended answer in the first case, but to return a SERVFAIL failure notice in response to the fragmented IPv6 response. When the local MTU on the server is lifted from 1,280 octets to 1,500 octets the Google resolver returns the server's DNS response in both cases.

It seemed curious to me at the time that Google's Public DNS behaved like this. It's not that this was unique to Google's service, by the way. As the article pointed out, the measurement experiment was able to identify many DNS resolvers that had similar problems in receiving fragmented IPv6 responses from an authoritative name server. However, the paradoxical problem for Google is that in all other respects their DNS service is incredibly good. It's accurate, it's fast, and it's honest in that it does not attempt to filter or manipulate what's in the DNS. So, in every other respect, Google's service just works really well. That means that the expectations that I have with Google's DNS service are high, and to see this lapse seemed all the more anomalous.

When I was preparing some material for a presentation on this problem at the November IETF meeting I had cause to check my original tests, and I was surprised to see that Google's behaviour was fixed! My attempts to feed Google fragmented IPv6 responses from the authoritative server all resulted in Google's DNS service accepting the fragments and correctly reassembling them. Something has changed for the better in the past few weeks! I have no idea how much work this entailed, but when you handle some 11%-12% of the entire Internet's DNS query load, then I can appreciate that making changes to the production service machinery requires some effort and a considerable level of care and attention to detail.

So, thanks Google! Thanks for making your public DNS resolvers work correctly in their handling of fragmented IPv6 responses. More generally I also want to say thanks to Google in demonstrating that they really care about getting IPv6 right. I appreciate it, and I'm sure many others do as well.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.